What is claimed is:

1. A speech recognition system for transforming an acoustic signal into a stream of phonetic estimates, comprising:

a frequency analyzer for receiving the acoustic signal and producing as an output a short-time frequency representation of the acoustic signal;

a novelty processor for receiving the short-time frequency representation of the acoustic signal, separating one or more background components of the representation from one or more region-of-interest components of the representation, and producing a novelty output including the region of interest components of the representation according to one or more novelty parameters;

an attention processor for receiving the novelty output and producing a gating signal as a predetermined function of the novelty output according to one or more attention parameters;

a coincidence processor for receiving the novelty output and the gating signal, and producing a coincidence output that includes co-occurrences between samples of the novelty output over time and frequency, wherein the coincidence output is selectively gated as a predetermined function of the gating signal, so as to produce a gated coincidence output according to one or more coincidence parameters; and,

a vector pattern recognizer and a probability processor for receiving the gated coincidence output and producing a phonetic estimate stream representative of acoustic signal.

2. A speech recognition system according to claim 1, wherein the short-time frequency representation of the audio signal includes a series of consecutive time instances, each consecutive pair separated by a sampling interval, and each of the time instances further includes a series of discrete Fourier transform (DFT) points, such that the short-time frequency representation of the audio signal includes a series of DFT points.

3. A speech recognition system according to claim 2, wherein for each DFT point, the novelty processor (i) calculates a first average value across a first predetermined frequency range and a first predetermined time span, (ii) calculates a second average value across a second predetermined frequency range and a second predetermined time span, and (iii) subtracts the second average value from the first average value so as to produce the novelty output point.

4. A speech recognition system according to claim 3, wherein the first frequency range, the first time span, the second frequency range and the second time span are each a function of one or more of the novelty parameters.

5. A speech recognition system according to claim 3, wherein the first predetermined frequency range is substantially centered about a frequency corresponding to DFT point, and the first predetermined time span is substantially centered about an instant in time corresponding to the DFT point.

6. A speech recognition system according to claim 3, wherein the first predetermined frequency range is substantially smaller than the second predetermined frequency range.

7. A speech recognition system according to claim 3, wherein the first predetermined time span is substantially smaller than the second predetermined time span.

8. A speech recognition system according to claim 3, wherein the second predetermined time span is large relative to the second predetermined frequency range.

9. A speech recognition system according to claim 3, wherein the second predetermined frequency range is large relative to the second predetermined time span.

10. A speech recognition system according to claim 3, wherein for each DFT point, the novelty processor further calculates one or more additional novelty outputs, and each additional novelty output is defined by characteristics including a distinct first frequency range, first time span, second frequency range and second time span, each characteristic being a function of one or more of the novelty parameters.

11. A speech recognition system according to claim 2, wherein the coincidence output includes a sum of products of novelty output points over two sets of novelty output points.

12. A speech recognition system according to claim 11, wherein the two sets of DFT points includes a first set of novelty output points corresponding to a first instant in time and a second set of novelty output points corresponding to a second time instance.

13. A speech recognition system according to claim 11, wherein the two sets of novelty output points all correspond to a single time instance.

14. A speech recognition system according to claim 11, wherein the coincidence processor performs the sum of products of novelty output points over two sets of novelty output points according to one or more selectably variable coincidence parameters including time duration, frequency extent, base time, base frequency, delta time, delta frequency, and combinations thereof.

15. A speech recognition system according to claim 2, wherein each of the time instances further includes an energy value in addition to the series of DFT points.

16. A speech recognition system according to claim 15, wherein the attention processor (i) compares the energy value to a predetermined threshold value according to a comparison criterion, so as to produce an energy threshold determination, and (ii) produces the gating signal as a predetermined function of the threshold determination.

17. A speech recognition system according to claim 16, wherein the one or more attention parameters include the predetermined threshold value, the comparison criterion and the predetermined function of the threshold determination.

18. A speech recognition system according to claim 1, wherein the novelty parameters, the attention parameters and the coincidence parameters are selected via a genetic algorithm.

19. A speech recognition system for transforming a short-time frequency representation of an acoustic signal into a stream of coincidence vectors, comprising:

a novelty processor for receiving the short-time frequency representation of the audio signal, separating one or more background components of the signal from one or more region of interest components of the signal, and producing a novelty output including the region of interest components of the signal according to one or more novelty parameters;

a coincidence processor for receiving the novelty output and the gating signal, and producing a coincidence vector that includes data describing co-occurrences between samples of the novelty output over time and frequency according to one or more coincidence parameters.

20. A speech recognition system according to claim 19, further including an attention processor for receiving the novelty output and producing a gating signal as a predetermined function of the novelty output according to one or more attention parameters, wherein the coincidence output is selectively gated as a predetermined function of the gating signal, so as to produce a gated coincidence output according to one or more coincidence parameters.

21 A speech recognition system according to claim 19, wherein the novelty parameters and the coincidence parameters are selected via a genetic algorithm.

22. A method of transforming an acoustic signal into a stream of phonetic estimates, comprising:

receiving the acoustic signal and producing a short-time frequency representation of the acoustic signal;

separating one or more background components of the representation from one or more region of interest components of the representation, and producing a novelty output including the region of interest components of the representation according to one or more novelty parameters;

producing a gating signal as a predetermined function of the novelty output according to one or more attention parameters;

producing a coincidence output that includes correlations between samples of the novelty output over time and frequency, wherein the coincidence output is selectively gated as a predetermined function of the gating signal, so as to produce a gated coincidence output according to one or more coincidence parameters; and,

producing a phonetic estimate stream representative of acoustic signal as a function of the gated coincidence output.

23. A method according to claim 22, further including (i) calculating a first average value across a first predetermined frequency range and a first predetermined time span, (ii) calculating a second average value across a second predetermined frequency range and a second predetermined time span, and (iii) subtracting the second average value from the first average value so as to produce the novelty output.

24. A method according to claim 22, further including calculating, for each of a plurality of DFT points from the a short-time frequency representation of the acoustic signal, one or more additional novelty outputs, wherein each additional novelty output is defined by characteristics including a distinct first frequency range, first time span, second frequency range and second time span, each characteristic being a function of one or more of the novelty parameters.

25. A method according to claim 24, further including performing a sum of products of novelty outputs over two sets of novelty outputs according to one or more selectably variable coincidence parameters including time duration, frequency extent, base time, base frequency, delta time, delta frequency, and combinations thereof.

26. A method according to claim 22, further including comparing the energy value to a predetermined threshold value according to a comparison criterion, so as to produce an energy threshold determination, and (ii) producing the gating signal as a predetermined function of the threshold determination.

27. A method according to claim 22, further including selecting the novelty parameters, the attention parameters and the coincidence parameters via a genetic algorithm.